
Information Systems January 11, 2019
Master of Science in Computer Engineering

Exercise 1 (9 points)
Let's consider the following relational schema:

PAINTING (PId, Subject, Painter, Technique, Support, Year, Price)
MUSEUM (Mid, Name, Director, City, Ticket)
EXIBITION (EId, PId,Title, MId, OpeningDate, ClosingDate)

Primary keys are underlined in the relations. Moreover, MId in EXIBITION is foreign key of MUSEUM, and PId in EXIBITION is foreign key of PAINTING. Techniques in PAINTING are Oil and Watercolour. Support in PAINTING is one of: wall, paper, canvas and wood. A painting appears in at least one exhibition. A museum appears in at least one exhibition.

Assume that:

$n_{PAINTING}$ =16.000	Val (Technique, PAINTING) = 2
n_{MUSEUM} = 125	Val (Support, PAINTING) = 4
$n_{EXIBITION}$ = 50.000	Val (City, MUSEUM) = 5

Given the query:
Title of exibitions in Parise with at least one oil painting painted on wood panel.

1) express the query as a relational-algebra expression;
2) show the basic steps of the query optimization process in terms of relational-algebra expression transformations
3) give an efficient strategy for computing the query.

Exercise 2 (6 points)
Consider the following schedule of concurrent transactions:
S: r3(z) r2(z) w1(x) r3(y) r1(z) w3(x) w2(y) w1(y) r1(x)

1) Show if S is conflict serializable (CSR) or view serializable (VSR). Explain why. If serializable, show equivalent serial schedules.

2) Apply the rigorous two-phase locking protocol to the schedule.
3) Apply the timestamp-ordering protocol to the schedule, assuming that aborted transactions are immediately restarted.

Exercise 3 (6 points)
Consider an empty B+-tree with m= 5.

1) Show the B+-tree after the insertion of the following values of the search key:
3 83 9 4 30 15 33 27 24 31 7

2) Show the form of the B+-tree after each operation of the sequence:
Insert 2; Insert 36; Delete 4.

Exercise 4 (9 points)

Let r=(A,B,C), with primary key A uniformly distributed on the interval [1; 10.000.000].

Assume
nr = 1.200.000 number of records in the relation
Lr = 100 byte size of a record (fixed length records)
LA = 8 byte size of attribute A
Lp = 4 byte size of a pointer
Lb = 1000 byte size of a block
Heap file organization on A.

1. Show (a) the minimum and (b) the maximum height of a B+tree tree index on search-key A.

2. Outline the steps in answering the following queries, the **best strategy** and the **cost** in terms of number of block transfers from disk in case (a):
1) select * from r where A=xxx;
2) select * from r where 100.000 <= A < 500.000;
3) select * from r where B=xxx